

EFFECTIVENESS ANALYSIS FOR WIDE AREA SEARCH MUNITIONS

David R. Jacques*, Air Force Research Laboratory, Munitions Directorate
101 W. Eglin Blvd. Suite 348, Eglin Air Force Base, Florida 32542
jacques@eglin.af.mil

Robert Leblanc*, SRS Technologies, ASI Division
838 N. Eglin Pkwy, Suite 202 Fort Walton Beach, Florida 32548
leblancr@eglin.af.mil

Abstract

The problem of estimating effectiveness of autonomous wide area search munitions will be considered. Due to the large potential search area of these munitions, terms such as single shot P_k are not adequate for evaluating expected success rates. Methods and models must take into account degradation due to false target attacks, and missed target acquisitions. This paper will establish some meaningful seeker metrics which are applicable to a wide variety of autonomous searching munitions. It will then establish a relationship between these metrics and analytical expressions for mission success. Although the analytic expressions will be limited to two specific scenarios, a discussion of their relation to more general Monte Carlo based effectiveness models will be presented. Finally, results for a specific scenario will be used to highlight an abbreviated seeker requirements flow down analysis for a wide area search munition.

I. Introduction

Autonomous wide area search munitions show great promise in being able to locate and engage widely dispersed and/or highly mobile or relocatable ground targets. They have the effect of decentralizing the search process from the strike aircraft or surveillance sensors to greater numbers of small, smart munitions with high resolution seekers operating at relatively short ranges. If equipped with Autonomous Target Acquisition (ATA) algorithms, these munitions can be delivered with very relaxed Target Location Error (TLE) requirements due to their ability to make autonomous target attack decisions. Hardware and ATA software development challenges receive most of the attention within the laboratory

environment, but there are also challenges associated with predicting the effectiveness of these munitions for a wide range of mission scenarios. The single shot P_k numbers associated with most direct attack munitions are not directly applicable to wide area search munitions because they do not account for the difficulty of searching over tens of square kilometers in order to find a target of interest. Terms such as False Alarm Rate (FAR), often associated with surveillance or search systems, are insufficient because they do not account for the effect of removing the munition from the battlefield if it fails to reject a clutter object or non-target vehicle. Further, FAR implementation in many current engagement level effectiveness models do not allow the seeker to reject the false target; by definition if it sees a false target, it engages it. For these reasons, meaningful metrics must be defined for wide area search munitions, and consistent procedures for incorporating these metrics into effectiveness models should be developed.

The problem being considered is that of a wide area search munition looking for a single target, while simultaneously trying to reject false targets that appear in the search area. Although this artificially limits the problem below the larger multiple target scenario we would like to address, it is important to fully understand how an effectiveness model handles the simple cases before we start introducing additional complexity to the overall problem. Said another way, how can we believe the model for the larger, more difficult cases if we can not reconcile it to what we know for these simple cases. Further, there are only minor differences in implementation between the single target scenario with multiple false targets, and a multiple target scenario. This paper will first establish some meaningful metrics which are applicable to wide area search munitions. The metrics will be defined and it will be shown how they apply to the entire class of wide area search munitions regardless of sensor type and degree of

* Major, USAF, Member AIAA

* Senior Analyst

This paper is declared a work of the U.S. Government and is not subject to copyright protection in the United States.

19981123 126

discrimination within the ATA algorithm. The paper will then discuss how these metrics can be incorporated into effectiveness models. Different implementation approaches will be shown, and significant differences in the results for the various approaches will be highlighted for certain cases. Where possible, the effectiveness model results will be compared to analytic solutions, and this will be used to suggest "correct" methods for modeling wide area search munitions. With some basic models in hand, effectiveness results and parametric trades for specific mission scenarios will be presented. These results will be used to highlight an abbreviated requirements flow down process for a wide area search munition whereby subsystem performance requirements are allocated to the seeker and ATA algorithms.

II. Evaluation Metrics

The problem being considered is that of a munition searching for a target or targets over an area significantly larger than a single scanned field of regard. We will define the Mission Set as a list of targets you wish to engage for a given mission, where destroying any target in the set constitutes a successful engagement. In order to find the target, the seeker has to prevent being fooled into thinking a false alarm once encountered is actually a real target. We need to draw a distinction between a false alarm (an object not in the Mission Set which triggers the ATA to evaluate an image or group of pixels) and a false target report (an object not in the Mission Set which is incorrectly confirmed as being a real target, thus causing an incorrect attack decision and loss of the munition). To highlight this distinction we will use the terminology of False Target Attack Rate ($FTAR$) to indicate the average rate ($/km^2$) at which munitions are expended on falsely confirmed targets. As will be shown in the following section, the $FTAR$ will be driven by the target you are searching for, the environment in which you are searching, and the type of seeker and ATA algorithm you are using to search.

The other key performance metric has to do with the ability to make a correct attack decision given that you encounter a real target. We will define a Target Report as a correct decision to attack a target in the Mission Set once it has been encountered. Note that we have not specified the degree of discrimination required on the part of the munition seeker. If the munition system relies on classification only, then a Target Report only requires correct classification. Further, if the munition relies on identification (ID) level of discrimination, the Target Report will require correct ID. Having said this, the metric of interest here is the Probability of Correct Target Report, or P_{TR} . It should be noted that while the definitions of $FTAR$ and P_{TR} do not reflect differing degrees of

discrimination, the quantitative values obtainable for these metrics will certainly reflect these differences. For instance, it is reasonable to expect that a system that relies on classification may obtain higher values for P_{TR} than one that relies on the higher level of discrimination associated with ID. However, the system with the identification level of discrimination will likely be able to achieve a lower level of $FTAR$.

The two metrics defined above, $FTAR$ and P_{TR} , can be considered Measures of Performance (MOP's) for an autonomous seeker. They do not represent a complete set of seeker effectiveness inputs; certainly parameters such as area rate of coverage and terminal guidance accuracy will also be important parameters for modeling and simulation. However, $FTAR$ and P_{TR} represent a minimal set of measurable seeker metrics that can be used to perform preliminary analytical assessments of system effectiveness as will be shown in the following sections.

At the system level, we will define Mission Success as achievement of the mission objectives. This could involve destruction of a single target, multiple elements within a site, or a specified percentage of targets within a larger formation. The system level metrics, or Measures of Effectiveness (MOE's), will be Probability of Mission Success (P_S), Mission Cost in terms of dollars or numbers of munitions, and ultimately we would like to be able to obtain an Estimated Kills per Sortie (EKS). We will not deal with EKS in this paper because it cannot be evaluated independently of the delivery platform. Further, it is relatively straight forward to take cost per kill in terms of numbers of munitions and map it to EKS once a loadout for a given delivery platform is known.

Figure 1 shows a simplified depiction of the relationship between MOP's and MOE's. Simply stated, this relationship is predominantly one of Modeling and Simulation (M&S). While Test and Evaluation (T&E) can be used to obtain or validate MOP's, we typically will not have sufficient test assets to fully evaluate the MOE's via hardware test. With increased use of Hardware-in-the-Loop simulations and synthetic imagery, it is also now possible to do some MOP evaluation via M&S. The remaining piece is one of mission planning which takes into account the various MOE's in deciding how many munitions to use and how to execute search and attack. It should be noted that many mission planning systems are now incorporating M&S to predict success rates and provide "optimal" mission plans. The remainder of this paper will concentrate on the M&S mapping between the MOP's and MOE's as shown in the top half of Figure 1.

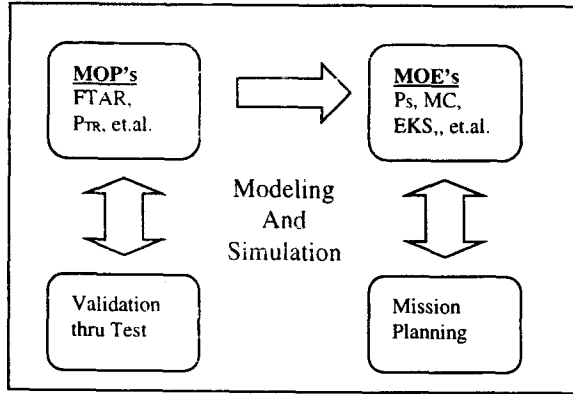


Figure 1. MOP/MOE Relationships




Figure 2. Incremental Search Area

III. Governing Equations - Single Munition Scenarios

Consider an autonomous munition looking for a single target in the area A_S of Figure 2, searching from bottom to top. The probability of successfully engaging a target in the incremental area ΔA is conditioned on not engaging a false target while searching area A . The incremental probability of success associated with killing a target in ΔA can be expressed as

$$\Delta P_S = P_{K|ENC} \cdot P_{ENC} \cdot \left(\frac{\Delta A}{A_S} \right) \cdot P_{FA}(A)$$

$P_{K|ENC}$ = Prob of kill given encounter

P_{ENC} = Prob that target appears in A_S

$P_{FA}(A)$ = Prob of not engaging a false target in A

An expression for P_S is obtainable by integrating $P_{FA}(A)$ over A_S , a closed form expression of which is obtainable as follows. Let

$\eta \equiv$ False target density (i.e. the density of non-target objects of suitable size/shape to potentially confuse the ATA)

$P_{FTAIFA} \equiv$ Conditional probability of false target attack given that a false target is encountered

$\alpha = FTAR \equiv$ False target attack rate (i.e. the product $\alpha = \eta \cdot P_{FTAIFA}$)

We will also need to define a false target attack probability distribution, $P(k, A)$. This represents the distribution of k , the number of false target attacks which would be reported

by the seeker in a non-commit mode, as a function of the area searched, A . It is a Poisson distribution with parameter

$$\lambda = \alpha \cdot A = \eta \cdot P_{FTAIFA} \cdot A$$

$P(k, A)$ can be expressed as

$$P(k, A) = \frac{[\eta \cdot P_{FTAIFA} \cdot A]^k e^{-\eta \cdot P_{FTAIFA} \cdot A}}{k!}$$

$$= \frac{[\alpha \cdot A]^k e^{-\alpha A}}{k!}$$

and the probability of searching A without executing a false target attack is

$$P_{FA}(A) = 1 - \sum_{k=1}^{\infty} P(k, A)$$

$$= P(0, A)$$

$$= e^{-\alpha \cdot A}$$

$$= e^{-\eta \cdot P_{FTAIFA} \cdot A}$$

Once we have this expression, we can now state our probability of mission success trying to engage a single target contained in a search area A_S as

$$P_S = P_{K|ENC} \cdot P_{ENC} \cdot \frac{1 - e^{-\alpha \cdot A_S}}{\alpha \cdot A_S}$$

The final term in this expression can be interpreted as the probability of having access to a single target within A_S by virtue of not having attacked a false target prior to encountering the real target.

The parameter P_{KIENC} contains all the probability factors associated with target acquisition, guidance and warhead lethality. It can be broken out as

$$P_{KIENC} = P_{LOS} \cdot P_{TR} \cdot P_{HTR} \cdot P_{KIH}$$

where P_{LOS} is the probability of having a clear line of sight to the target (statistically dependent on the terrain and the seeker depression angle), P_{HTR} is the probability of hitting the target given that you found it and correctly recognized it (primarily a guidance parameter), and P_{KIH} is the probability of killing the target given that you hit it (a warhead parameter). Although not shown, the equations presented here can be generalized to accommodate reliability factors, and this typically is done for most effectiveness models.

The decomposition of the $FTAR$ into a product of false target density (η) and the probability of false target attack given a false alarm (P_{FTAIFA}) is a point that warrants further discussion. For this breakout of $FTAR$, we should note that η is completely determined by what you are looking for, where you are looking to find it, and the nature of the seeker (e.g. millimeter wave vs. imaging infrared vs. LADAR). Once the basic seeker design has been chosen, η is scenario dependent. On the other hand, P_{FTAIFA} is determined by how good the seeker and ATA are in terms of rejecting false targets. Finally, because η is scenario dependent, we should expect significant variation in the $FTAR$ as we vary the mission targets and clutter background. For example, we should obviously expect a much higher value of both η and the $FTAR$ if we are looking for a command van in an urban environment as opposed to a large mobile missile launcher in a desert environment.

We have said that we can view $FTAR$ in terms of two components, but it is not true that only the product of the two numbers is important to the formulation. For instance, consider the case of a 1x100 km strip being searched by a munition with a 1 km wide seeker footprint. Further, let us assume that the system $FTAR$ is 0.01 for this particular scenario. One approach to formulating the effectiveness problem, albeit flawed as shown in the following example, is to interpret $FTAR$ as a false target density, and to randomly place a single false target and a single real target in the 100 km² search area (for a false target density of 0.01). This method typically specifies P_{TR} , P_{HTR} , and P_{KIH} for the real target if it is encountered, but the probability of attacking the false target is identically 1 if encountered. Note that this formulation satisfies the necessary condition for the $FTAR$, because $P_{FTAIFA} = 1$ and $\alpha = \eta \cdot P_{FTAIFA}$. While this may appear correct, consider

the case where there is no real target in this search path. The probability of the munition falling for the false target is identically 1, because a false target density of 0.01 will always result in one false target being placed in the 100 km² area, and the munition will always encounter it and always take itself out of the game by attacking it. Realistically speaking however, a $FTAR$ of 0.01 does not say that, with probability 1, you will falsely attack a non-target while searching a 100 km² area, even if there are no real targets in the area. There is a non-zero probability of not false alarming while covering the search area that is not accounted for when you assume

$$\alpha = \eta$$

$$P_{FTAIFA} = 1.0$$

What the $FTAR$ should represent is your ability to reject false alarms as they continually "bombard" the seeker and ATA. This suggests that a better approach may be to put greater numbers of false targets in the search area, but specify some ability on the part of the seeker and ATA to reject the false targets. The product of η and P_{FTAIFA} needs to be held constant and equal to the desired $FTAR$, but in the limit (as η gets large) this type of formulation should converge to the "true" solution.

The next question we might ask is whether or not this discrepancy is "in the noise" numerically, or does it warrant consideration for alternate formulations. In fact, the differences between these two approaches is not only significant, but far exceeds the difference other system trades can make in the overall effectiveness calculations. For instance, consider a single munition with

$$P_{KIENC} = 0.72$$

$$A_s = 50 \text{ km}^2$$

Figure 3 shows a comparison between two different approaches to the effectiveness calculation. The "Ana" line shows the analytic solution as described above. The "MC1" line shows a representative Monte Carlo effectiveness model. MC1 uses a False Alarm Rate (FAR) parameter to specify a false target density, and sets $P_{FTAIFA} = 1$. Note the difference in the results for the two approaches as $FTAR$ increases above 0.01. The third curve in the chart (MC2) refers to a modification of the MC1 model to handle false targets in a different way. What we did was specify that there were 50 non-targets in the search area ($\eta = 1$ in this case) and we specified a probability of correctly identifying these non-targets as non-targets. The probability of correct ID (P_{id}) for the non targets was specified as

$$P_{id_{NT}} = 1 - P_{FTA/FA}$$

$$\alpha = \eta \cdot P_{FTA/FA}$$

While this is still just an approximation to the analytic solution, the curve shows that it is a much closer approximation than the previous implementation.




Figure 3. Single Munition Success vs. FTAR

IV. Extensions for Multi-Munition Engagements

The result discussed above has an even greater impact to mission effectiveness calculations when multiple munitions are used to achieve higher mission success rates. In addition to the problem of how to distribute false targets around the search area, we now need to consider whether or not we have correlated munition behavior at the real or false targets. What we are referring to is whether or not the behavior of one munition at a potential target (real or false) affects the behavior of subsequent munitions which encounter the same target. For several munitions flying over the same search path, in the same direction, with the same mission set loaded, you would expect a high degree of correlation in how the munitions behave when encountering a stationary target. They will see the targets at the same range and aspect angle, so it is reasonable to assume that they will make similar decisions. However, if the munitions see the targets (real or false) at different ranges or aspect angles, it is not nearly as clear as to what the degree of correlation should be between the munitions. For very different aspect angles it could be argued that you should expect a relatively small degree of correlation between the false target behavior of the munitions. It should be noted that, for the case of more potential false targets than real targets (not unrealistic for the Lethal SEAD and Theater Missile Defense/Attack Operations missions), the case of correlated behavior will

result in lower mission success rates than the uncorrelated case for the same number of munitions.

If an effectiveness model performs a random draw each time a target is encountered, this essentially assumes uncorrelated behavior because there is no change to the probabilities based on previous munition behavior. With this, we note that the MC1 model assumes uncorrelated behavior at the target (separate draws), but 100% correlation at a false alarm ($P_{FTA/FA} = 1$ states that all munitions encountering a given false target will attack it with probability 1). A contributing factor for this is that the positions of all targets (real and false) is set at the beginning of each trial and remains constant for all munitions encountering it. Having said this, it is apparent that this approach to the problem is the most conservative (or pessimistic) method of evaluating mission success rates. The opposite situation where you have correlated behavior at a real target and uncorrelated behavior at a false target should yield a much higher mission success rate calculation.

To partially illustrate the argument from the previous paragraph, Figure 4 will show a simple comparison for the same case as discussed above, but now with several munitions overflying the same corridor. The analytic solution for multiple munitions overflying the same path, assuming uncorrelated behavior at both false and real targets, is given by

$$P_S^N = \frac{P_{ENC}}{A_s} \sum_{n=1}^N \left[\left(1 - (1 - P_{K|ENC})^n \right) \cdot I_n^N \right]$$

$$I_n^N = \frac{N!}{n!(N-n)!} \int_0^{A_s} \left(e^{-\alpha \cdot A} \right)^n \left(1 - e^{-\alpha \cdot A} \right)^{N-n} dA$$

A similar expression can be derived for the case of opposing paths, but the most general case of arbitrary path angles is complicated enough that we will not derive it analytically. A properly constructed model can easily evaluate this more general case numerically. Further, the simple case of overlying paths is sufficient to make the point. The "Ana" line depicts the analytic solution for two munitions looking for a single target along the same path with uncorrelated behavior at both real and false targets. The "MC1" line shows the Monte Carlo results for this case. Once again, the MC1 model uses a specified FTAR as a false target density, thereby randomly distributing these false targets across the search area and assigning a $P_{FTA/FA}$ of 1.0 given that one of the false targets is encountered. This is the very conservative case mentioned above which assumes uncorrelated behavior at a real target and 100% correlated behavior at a false target. If the same modifications as mentioned above for the single munition case are made for the case of two munitions, than the model

can predict the analytic solution (the curve noted as "MC2"). Specifically, we have specified

$$P_{idNT} = 1 - P_{FTA/FA}$$

$$\alpha = \eta \cdot P_{FTA/FA}$$

and using 50 targets in a 50 km² area results in $\eta = 1$. The final curve uses the same modification, but increases the false alarm density to $\eta = 2$ (while holding $FTAR$ constant). This curve shows even closer convergence to the analytic solution. The drawback to using a higher false alarm density is that the simulation slows down significantly in order to handle increasingly higher numbers of target "interrogations". While this method is not necessarily "the correct" solution, it does treat all targets (both real and false) equally, and it better represents what an $FTAR$ figure should mean in terms of mission effectiveness.




Figure 4. Two Munition Success vs. FTAR

It is also instructive to compare the multi-munition equation above with a simplified roll-up approach for a multi-munition engagement. The simplified roll-up for N munitions is given by

$$\hat{P}_S^N = P_{ENC} \cdot (1 - (1 - P_S^1)^N)$$

$$P_S^1 = P_{K|ENC} \cdot \frac{1 - e^{-\alpha \cdot A_S}}{\alpha \cdot A_S}$$

This expression is not truly valid for our wide area search formulation because it assumes a different random target placement for each munition. In reality, all munitions will be looking for the same target or targets which can not move instantaneously. We can compare this graphically with the multi-munition equation presented earlier, and the results are shown in Figure 5. The curve labeled "Ps1" represents the analytic solution for two munitions overflying the same path, while "Ps2" represents the analytic solution for two munitions overflying the same

area via opposing paths. The final curve, "Ps3", represents the simple multi-munition roll-up approximation. Figure 5 shows that the simple multi-munition roll-up is optimistic when compared to the same path formulation (Ps1), but not when compared to the opposing path approach (Ps2). Once again, the assumption of uncorrelated behavior (at either a real or false target) is not strictly valid, and we should expect a high degree of correlation for the case where the munitions are traversing the same path in the same direction. For scenarios where the potential false targets greatly outnumber the real targets, correlated behavior will degrade the overall mission success rate. For this reason, search patterns should be planned which decrease the degree of correlated behavior at false targets. This can be done through the use of lateral offsets between munitions and/or different approach vectors. While this does not make the assumption of uncorrelated behavior valid, it reduces the error associated with this assumption.




Figure 5. Comparison of Analytic Formulations

V. Seeker Requirements Flow Down

Armed with an analysis approach, we can now perform a seeker requirements flow down based on a system level effectiveness analysis. The scenario we will use for an example is a single mobile missile launcher target. The initial target location error is assumed to be 1.25 km, and we assume the target can move in any direction at a rate of 25 km/hr beginning 5 minutes after a launch event. Further, we assume no target updates after the launch event. The delivery platform for the wide area search munition is located 90 km away from the target, and either the munition or the launch platform can ingress to the target at Mach 0.9. The ingress can begin no sooner than 5 minutes after the target launch event for this analysis. Once the munitions are in their search patterns, they have a linear area rate of coverage equal to 4.5 km²/min. Actual coverage rate is dependent on search altitude, velocity, scan width, and pixel generation and processing rates. Likewise, the velocity and scan width must be consistent with the

munition maneuverability to ensure targets at the edge of the scan can be engaged. Although all of these factors and more must be considered in a thorough seeker requirements flow down, we will limit the discussion to the ATA performance metrics introduced earlier, namely $FTAR$ and P_{TR} .

Figure 6 shows the growth of the Target Location Error (TLE) and the number of munitions required to contain the uncertainty growth. With no knowledge of which direction the target is likely to run, the uncertainty grows with the square of time, and the uncertainty due to target motion quickly becomes the dominant factor. It is obvious from this chart that a single munition is not sufficient to contain this worst case uncertainty growth. Either additional munitions will be required, or the target search must take into account Intelligence Preparation of the Battlefield (IPB) to determine likely egress routes for the target vehicle. Assuming no IPB and no overlapping coverage, it will take at least 4 munitions to contain the uncertainty growth, with each munition covering 12-13 km^2 . The total area which needs to be covered for the 4 munition case is approximately 50 km^2 . Alternatively, 8 munitions could be used, with each munition covering approximately 4 km^2 and a total search area required of 30-35 km^2 . This analysis can easily be repeated for other mission-target scenarios, and it can be used to provide an indication as to the total search area required and the ballpark amount each munition needs to be able to search. However, this type of geometric analysis can only provide indicators because it ignores munition reliability factors, false alarm behavior, and missed target probabilities among others. The analytic results presented earlier can be used for effectiveness results for simplified cases, and simple Monte Carlo models (if properly implemented) can provide meaningful results for more general cases.

Figure 7 illustrates the extreme dependence of mission success and the number of munitions to achieve success on the $FTAR$. We have assumed $P_{KIENC} = 0.72$, which may be optimistic, but will serve to illustrate the important trends. Curves such as this illustrate where "the knee in the curve" is, beyond which a prohibitively high number of munitions will be required to achieve a given level of mission success. A projected mission cost can be obtained from this curve by reading the number of munitions required for a desired confidence level, multiplying by the average unit procurement price of the munition, and adding in any additional costs associated with the delivery platform sortie cost. It should be noted that Figure 6 illustrated that multiple munitions may be required to contain the uncertainty growth, and Figure 7 adds to this the fact that multiple munitions will be required to overcome all but the most optimistic False Target Attack Rates. Certainly this calls for both a low value for the

$FTAR$ and a low cost for the wide area search munition. It also suggests that IPB should be considered wherever possible in mission planning to reduce the number of munitions required, or increase the P_S .




Figure 6. Munition #'s vs. Growth in TLE




Figure 7. Munition #'s vs. FTAR and Mission Success

Figure 8 once again shows the sensitivity to $FTAR$, but we now parameterized around P_{TR} for a given level of mission success required. It is interesting to note that the single fleeing target scenario is relatively insensitive to P_{TR} . This suggests that, for this mission, P_{TR} could be traded off in the seeker operating characteristics in order to obtain a lower level of $FTAR$ (for a given seeker and algorithm, the two parameters will be at odds with each other). This also suggests that a higher level of overall discrimination will likely be required for the wide area search munition.




Figure 8. Munition #'s vs. FTAR and Target Report

VI. Conclusions

Wide area search munitions have the potential for revolutionizing the deep strike and interdiction mission areas. While there are significant challenges in the component technologies to realize this potential, there are equal challenges associated with optimum employment and effectiveness analysis for these types of high leverage munitions. The added value obtained by an autonomous wide area search munition is not currently modeled at the engagement, mission, or campaign level. Wide area search capability by low cost autonomous munitions provides a significant force multiplier. It allows area sweeps by swarms of munitions to perform preemptive destruction of enemy air defenses and to locate and destroy high value surface-to-surface mobile missile launchers. The introduction of these munitions will require a paradigm shift in the way we currently model munitions in order to show their military worth at the campaign level. This paper provides an initial examination of a subset of performance metrics that can be incorporated into engagement level effectiveness models and simulations to capture the significant capability these munitions offer the theater commander to locate and destroy time critical targets.

Acknowledgements

The authors would like to acknowledge and thank several individuals who contributed directly and indirectly to this paper. First, Dr Tom Davis of the Munitions Directorate of Air Force Research Lab was instrumental in providing meaningful and measurable seeker metrics as part of his work with the Anti-Materiel Munition Integrated Product Team. He also provided an invaluable sanity check as we progressed in our analysis. We would also like

to acknowledge the teams from Lockheed Martin Vought Systems and Raytheon TI Systems who performed detailed seeker requirements flow down studies for an autonomous wide area search munition. In particular, we enjoyed countless hours of discussion with Dr Lewis Minor, Carl Cowan, and Ar Gant from Lockheed Martin, and Dr Bob Henderson and Janice Huston from Raytheon. Their knowledge and enthusiasm for this subject area made it a genuine pleasure to work with them.